

PGPUB-DOCUMENT-NUMBER: 20020114326

PGPUB-FILING-TYPE: new

DOCUMENT-IDENTIFIER: US 20020114326 A1

TITLE: Address mapping mechanism enabling multi-domain
addressing in communication networks

PUBLICATION-DATE: August 22, 2002

INVENTOR-INFORMATION:

NAME	CITY	STATE	COUNTRY	RULE-
47				
Mahalingaiah, Rupaka	Austin	TX	US	

US-CL-CURRENT: 370/389, 370/386

ABSTRACT:

An architecture, system and method are provided for transparently mapping addresses across multiple addressing domains and/or protocols. A destination of a packet can therefore be transferred from a first addressing domain within one network to a second addressing domain within another network, without inserting knowledge into the packet of the relationship between the two separate and independent domains. Transmission modules within one network can be identified with unique identification numbers or addresses assigned during configuration of those modules. The identification numbers assigned internal to the network can be mapped and placed upon the packet as the packet enters the network. Mapping, however, is minimal, knowing that relatively few external devices are connected to select internal devices and/or modules. The packet can then be mapped into the network, where it is then transferred across the network whereupon it is mapped to another network or termination device external to the network. The downstream network or termination device may have an addressing domain entirely separate and independent from the addressing domain used within the network, the benefit of which is to free up external addressing space that would normally be used if the external addressing domain were used internally.

----- KWIC -----

Abstract Paragraph - ABTX (1):

An architecture, system and method are provided for transparently mapping addresses across multiple addressing domains and/or protocols. A destination of a packet can therefore be transferred from a first addressing domain within one network to a second addressing domain within another network, without inserting knowledge into the packet of the relationship between the two separate and independent domains. Transmission modules within one network can be identified with unique identification numbers or addresses assigned during configuration of those modules. The identification numbers assigned internal to the network can be mapped and placed upon the packet as the packet enters the network. Mapping, however, is minimal, knowing that relatively few external devices are connected to select internal devices and/or modules. The packet can then be mapped into the network, where it is then transferred across the network whereupon it is mapped to another network or termination device external to the network. The downstream network or termination device may have an addressing domain entirely separate and independent from the addressing domain used within the network, the benefit of which is to free up external

addressing space that would normally be used if the external addressing domain were used internally.

Title - TTL (1):

Address mapping mechanism enabling multi-domain addressing in communication networks

Summary of Invention Paragraph - BSTX (3):

[0002] This invention relates to a communication system ("network") that transparently maps and operates in one domain of addressing that may be different from the addressing domain interfacing with the network. More specifically, the invention provides mapping between an addressing domain outside the network with an addressing domain inside the network at end nodes of the network, the latter addressing domain not being constrained to a particular bit size and being used to supplant addresses outside the network that, but for this invention, would be used inside the network as well.

Summary of Invention Paragraph - BSTX (16):

[0015] The mapping of a domain of addresses should ideally be transparent to the division of a global network into possibly numerous local networks. An improvement would exist if somehow the local networks (or "networks") could have different addressing domains. A network being serviced could beneficially have an addressing domain and/or protocol that is separate and independent from an addressing domain of a network providing the service. This would enable efficient organization and management of data transmission within the domain providing the service without impacting the address domain that is being serviced. Any modifications and scaling within the network providing services will then be ideally isolated from the network being serviced. Similarly, any changes or scaling to the network being serviced will be transparent to the servicing network. The aforementioned concepts, albeit not available in conventional addressing or address mapping would, if employed, significantly enhance such conventional systems.

Summary of Invention Paragraph - BSTX (20):

[0018] The problems outlined above are in large part solved by an addressing mechanism that allows transparent communication between multiple addressing domains. For example, a servicing network may have an addressing domain separate from a network being serviced. The servicing network can hereinafter be regarded as the "network" or "networks" while the serviced network is external to the servicing network. Contained within the servicing network are intermediate nodes placed between end nodes. The nodes are hereinafter referred to as "modules." The modules perform routing functions based on source and destination addresses of the packet being forwarded. It is, therefore, contemplated that the communication system hereof includes one or more networks interconnected with one another between termination devices. A termination device may be the source of data, and another termination device may be the destination for that data. The termination devices can be computers, routers, gateways, switches, bridges, or any hardware or software module which can transfer, map, read data into the network, or write data from the network. Reading and writing functions can be those normally associated with any media usable by a hardware or software element, and the media can be either permanent or temporary.

Summary of Invention Paragraph - BSTX (25):

[0023] In order to transfer a packet of data from a termination device to

another termination device, separated by a network with an addressing domain internal to the network different from addressing domains external to the network, a simple mapping function is all that is required. In other words, the termination device which sources the data may have a unique Ethernet address domain noted on the packet being sent. Moreover, the destination termination device has its own unique addressing domain shared with that of the source termination device. As such, the source and destination devices are said to have a common addressing domain, where each device has its own unique identifying number. In addition to the termination devices having a unique identifying number of a first addressing domain, the entry and exit end modules also have a unique identifying number within the first addressing domain. In this fashion, the source termination device will have an identifying number separate and distinct from the identifying number of the entry end module. In addition, the exit end module will have an identifying number unique to the exit end module within the same addressing domain as the entry end module, the source termination device and the destination termination device. In this fashion, a single addressing domain can be used to forward data from the source termination device to a unique entry end module and eventually to an exit end module coupled to the destination termination device. What happens within the network, and among intermediate modules, is entirely separate and independent of addressing between termination devices and end modules. In other words, the intermediate modules can have an addressing scheme altogether different from the addressing scheme used to send data into an entry end module and receive data from an exit end module. The routing between termination devices and end modules is, therefore, said to be independent of the routing between intermediate modules. Independent routing is often referred to interchangeably as independent addressing. By freeing addresses used in Ethernet from the intermediate modules, it will require considerably less number of Ethernet addresses only on the periphery (end module) and termination device connected thereto. Therefore, an extremely large network can be easily accommodated within the existing 48-bit Ethernet address range or the 32/64-bit IP address range.

Summary of Invention Paragraph - BSTX (31):

[0029] According to another embodiment, the communication network can include a pair of end modules. The pair of end modules, known as entry end modules and exit end modules operate as bookends between intermediate modules. The first addressing domain is used to identify each of the end modules and the intermediate modules according to unique identification number assigned to such modules. The second addressing domain is used to identify each of the end modules and all modules external to the network with unique identification numbers. The identification numbers of modules within the first addressing domain can overlap with identification numbers of modules within the second addressing domain. However, since the first and second addressing domains are separate and independent from one another, there is no chance that addressing between modules of the first addressing domain can be improperly routed to modules within the second addressing domain. A mapping occurs between an identification number of an end module represented within the first addressing domain to a corresponding identification number within a second addressing domain and vice-versa as data is transferred into and out of the end modules. By using dissimilar addressing domains, and converting or mapping identification numbers in the two domains at the end modules, a substantial amount of public domain identification numbers of the external second addressing domain are not needed internal to the network.

Brief Description of Drawings Paragraph - DRTX (8):

[0037] FIG. 6 is a block diagram showing an example by which Ethernet addresses and IP addresses are assigned to modules outside the first addressing

domain, with addresses assigned to modules within the network (i.e., end module and intermediate modules), wherein a mapping occurs at the end modules between first and second addressing domains;

Detail Description Paragraph - DETX (3):

[0044] Each node or module 12 may embody a subnet or a plurality of interconnected subnets. Select modules 12a and 12b can be used to receive input data into network 10, or transmit output data from network 10. Contained within network 10 can be intermediate modules 12c, 12d, and 12e. Intermediate modules 12c-12e are used solely to transfer data within the network, but do not operate as an end node for transmitting data external to the network. Network 10 may be one of possibly numerous networks interconnected with one another to form an internet 15. Thus, an internet is one which forwards data between termination devices 14. The termination devices can operate as bookends between one or more networks 10. Shown in FIG. 1 are termination devices 14a and 14b which transmit information across network 10. End modules 12a, 12b, and/or 12c can be connected to end modules of another network, or to termination devices 14. In the embodiment shown in FIG. 1, three networks are interconnected with one another, and labeled as network 2 (NW.sub.2), network 3 (NW.sub.3), and network 4 (NW.sub.4). For purposes of further illustration, network 10 is used to illustrate mapping between addressing domains internal and external to network 10 as set forth below.

Detail Description Paragraph - DETX (12):

[0053] FIG. 4 illustrates traffic manager 22 shown in more detail. Traffic manager 22 functions to direct incoming data to an appropriate output port. It determines the output port for an incoming data packet based on a simple algorithm. The traffic manager is aware of its position relative to the network and its neighboring traffic managers or nodes. Based on this knowledge, the traffic manager determines the appropriate output port in accordance with the predetermined algorithm. The traffic manager 22 also interfaces between ports of different media, speed, and direction. After the output port assignment is made, the traffic manager can (i) store the data in a local buffer if the incoming data is slower than the speed of the directed-to output port until there is sufficient time to sustain the speed of the output port, (ii) forward the data without an intervening storage if the output port is free or if the incoming data is at the same speed as the output port, or (iii) transfer the data through a FIFO-type storage device when the incoming data is coming at a faster rate than the speed of the output port. The traffic manager also manages resource allocation, such as allocation of various output ports and buffers, and the algorithms make decisions regarding allocation to reduce the latency through the traffic manager.

Detail Description Paragraph - DETX (19):

[0060] Transfer of a wrapped packet from an originating termination device 14a to a destination termination device 14b entails numerous mechanisms. Reference numeral 37 illustrates a wrapped packet ensuing from entry module 16a. The wrapped packet includes a source address (SA) of an originating termination device having a unique identification number, and also includes a destination address (DA) of the exit module (MOD2 ADDR) 16b. The source address is therefore the address of termination device 14a which gets translated (or mapped) to an address of entry module 16a. The source address of termination 14a is mapped to an address of entry module 16a by noting the relationship between the addressing domain external to NW.sub.2. That addressing domain may be, for example, an Ethernet addressing domain of the second addressing domain D.sub.1. However, it is important to change, modify, or map the addressing domain of the module 16a of the Ethernet address to an

addressing domain unique and internal to NW.sub.2, that addressing domain being the first addressing domain D.sub.2. Mapping between the first and second addressing domains at the entry end module proved beneficial when the destination termination device replies to the originating termination device.

Detail Description Paragraph - DETX (20):

[0061] Upon receipt of the wrapped packet by exit module 16b, the destination address changes (i.e., is mapped) from exit module 16b to an addressing domain recognizable to destination termination device 14b. In this manner, mapping occurs when entering the network and also occurs when exiting the network. Additionally, mapping is needed to indicate possibly numerous termination devices attached to module 16b. Thus, mapping is used to translate between addressing domains of the entry and exit modules and intermediate modules of NW.sub.2 and the addresses of the entry and end modules which then have to forward the data to corresponding termination devices. The source address, however, remains the same as that which was changed when the packet is forwarded into the transmission network. Given that the destination address is termination device 14b, exit module 16b performs the entirety of all mapping needed by the transmission network. Mapping occurs in order to determine which of the possible N number of termination devices 14b is to receive the outgoing packet.

Detail Description Paragraph - DETX (21):

[0062] Mapping proceeds by comparing the address within the packet against the appropriate identification number of a possible N number of termination devices T.sub.1/T.sub.2. A mapping table or a generalized selection unit may therefore be formed within memory of exit module 16b. Mapping is performed by comparing the destination address within the wrapped packet against the identification number of a termination device selectibly coupled to the exit module 16b. When the comparison yields a match, the packet is forwarded to the appropriate termination device.

Detail Description Paragraph - DETX (23):

[0064] As noted by reference numeral 37, the wrapped packet contains source and destination addresses, whereby the source address of an external terminal identification number/address is mapped to an internal module identification number/address residing at an address domain separate and independent from the external identification number/ address. Moreover, the destination address of the wrapped packet within the network has an identification number/address unique to the internal identification numbers of the network. Accordingly, the internal addressing domain/protocol is shown with the letter "D.sub.2" as a first addressing domain of unique identifier numbers for that network (intermediate modules and end modules) which may share with numbers/addresses outside the network, however, addressing outside the network is separate and distinct from addressing internal to the network by virtue of a mapping function which occurs at the periphery of the different addressing domains (whether an addressing domain can be encompassed in one, two, or more networks). The internal addressing domain D.sub.2 and the external addressing domain D.sub.1 are consistent with those shown in FIG. 1.

Detail Description Paragraph - DETX (24):

[0065] FIG. 6 illustrates the transparent conversion and/or mapping of addresses of domain D.sub.1 to domain D.sub.2 addresses external to the network. Certain numbers are applied to those addresses, to assist in understanding the mapping operation. It is understood, however, that the addressing numbers are merely examples, and in no way limit the scope and

breadth of the mapping operation. Termination devices are resident in addressing domain 1 (D.sub.1 of FIG. 1). The end modules (such as 54 and 56) reside both in serviced domain 1 and servicing domain 2 (D.sub.2 of FIG. 1). End modules are responsible for transparently mapping the addresses of the two domains. Intermediate modules (such as 50 and 52) are resident entirely in D.sub.2.

Detail Description Paragraph - DETX (25):

[0066] FIG. 6 depicts a scenario as applied to popular IP Networking layer and Ethernet Data Link Layer or Physical layer. It should be noted that the transparent address mapping mechanism applies to other transmission protocols as well. In the example shown, termination devices can be regarded as host devices and, in the domain D.sub.1 (Ethernet) addressing, can have numerous clients coupled to that host, each having its own address. The host device can have a domain D.sub.1 Ethernet address "EA" unique to that host. Information can be sent to the host or received from the host by clients coupled thereto, each client having its own unique address represented at a higher OSI layer (e.g., an internet protocol address "IP"). Thus, for example, IP 10 can represent the IP address of a termination device which forwards information into the transmission network. Since device at address IP 10 is presented in the Ethernet physical layer, source address IP 10 is wrapped with another source Ethernet address EA 1008. The source IP address and domain D.sub.1 source Ethernet address can be noted as "SIP" and "SEA" shown in FIG. 7.

Detail Description Paragraph - DETX (28):

[0069] In the example by which data flows as packets from IP 10 to IP 100, end module 54 suffices as an entry end module for initial transmission, and end module 56 functions as an exit end module for that transmission. During a subsequent reply, the entry end modules are reversed, whereby end module 56 suffices as an entry end module and end module 54 will suffice as an exit end module. Thus, the end modules function as bookends for data transfer across the network, whereby the entry end module wraps the packet and domain D.sub.1 source and destination addresses with domain D.sub.2 source and destination addresses. The exit end module strips the domain D.sub.2 source and destination addresses, thereby leaving in place mapping to the final target device, represented within the network as an Ethernet device, or outside the domain D.sub.2 network as an IP-identified device.

Detail Description Paragraph - DETX (29):

[0070] Turning to FIGS. 6 and 7, in conjunction, illustration of an ARP process is shown. The ARP is shown with an example by which the relationship between end modules and modules/devices external to the domain D.sub.2 network are resolved. Specifically, the relationship between the source D.sub.2ID ("SD.sub.2ID") and the devices external to the network coupled to the SD.sub.2ID are resolved, along with the relationship between the destination D.sub.2ID ("DD.sub.2ID") and the devices external to the network. The broadcast operation is better illustrated by the example of which SIP at address 20 forwards a broadcast signal (SEA 1008) to all modules and termination devices within the transmission network as well as devices connected to that network to determine the D.sub.2ID of an exit end module coupled to the DIP of 100, not yet knowing the DEA of the device corresponding to the DIP. As such, a broadcast signal of all ones, can be sent across the network from the SIP to the DIP. This entails wrapping the SIP, DIP, SEA, and DEA with the SD.sub.2ID and DD.sub.2ID. However, the DD.sub.2ID is not yet known. The SD.sub.2ID will be the D.sub.2ID of the entry module (54 in FIG. 5). Entry module also has the SEA value of 2010 of domain D.sub.1. Domain D.sub.1 SEA (1008 in the example) of incoming packets is mapped to the domain

D.sub.1 SEA (2010) of the entry module. If there are N termination devices attached to given entry module, this mapping will be a N:1 mapping. Moreover, the SEA-mapped address is translated or mapped on a 1:1 basis to the SD.sub.2ID.

Detail Description Paragraph - DETX (30):

[0071] The mapped packet of data is then received by the exit end module, whereby the SD.sub.2ID and DD.sub.2ID are stripped from that wrapped packet. The DEA remains an unknown, hence the packet is sent to all of the termination devices as it is a broadcast packet. When a terminating device responds, when forwarding the broadcast response back to the originating (source) device, the DIP is mapped to the SEA. Similar to when the broadcast ARP is wrapped during the initial transmission, the SEA is mapped to the entry end module of the responding transmission. As shown, the SEA of AB12 (hex) is mapped N:1 to the entry end module at 2002 (hex). Thus, the entry end module of the response (i.e., module 56) can transfer the wrapped packet from source Ethernet address SEA AB12 back to IP 20. As shown, the response packet will have a SIP of 100, DIP 20, SEA of AB12 and DEA of 2010 when entering the module 56. The DD.sub.2ID is determined by using a static 1:1 mapping table of domain D.sub.1 DEA to domain D.sub.2 D.sub.2ID. The packet will be wrapped with SD.sub.2ID of 106 and DD.sub.2ID of 100. The termination device SEA of AB12 will be replaced with the entry module SEA of 2002.

Detail Description Paragraph - DETX (31):

[0072] At the exit end module 54, the internal structured identification numbers (SD.sub.2ID and DD.sub.2ID) are removed, as shown. Also, the DEA of the exit termination device is determined by using the terminating device DIP as the look-up parameter for 1:N mapping to domain D.sub.1 terminating device DEA. The end module domain D.sub.1 DEA (2010) is replaced with DEA (1008) of the terminating device corresponding to the DIP of 20. While the relationship between the entry termination device and the entry end module is known by virtue of the packet entering the entry end module and getting wrapped with the entry end module identification number, the ARP operation beneficially discovers the relationship between the exit end module and the targeted, or downstream, device. By broadcasting packets of data with a destination noted as a known broadcast destination, response to that broadcast operation will determine not only the relationship between domain D.sub.1 Ethernet address and the IP address of the targeted device, but also the relationship between the targeted device adjacent to the network and the exit end module. Thereafter, normal data transmission can occur to a DEA of the targeted device as well as the DD.sub.2ID of the exit end module adjacent that targeted device.

Detail Description Paragraph - DETX (32):

[0073] In one embodiment, the mapping of DIP to DEA can be entered at the time in which the network is configured. The mapping is generally configured once (i.e., statically); however, the network can be updated to reflect any changes to the transmission network. In another embodiment, The mapping can be captured every time a packet is passes through the entry module. As can be seen, there can be various embodiment with combinations of these mechanisms.

Detail Description Paragraph - DETX (33):

[0074] Thus, by using simple 1:1, N:1, and 1:N mappings, transparent address mapping between multiple domains is achieved for data transmission. All the mapping occurs at the boundary resulting in a simple, transparent address mapping mechanism.

Detail Description Paragraph - DETX (34):

[0075] FIG. 8 illustrates in further detail a data transmission operation, knowing the DD.sub.2ID to be adjacent to the DIP. In other words, data transmission following ARP allows for the entry end module to wrap the packet of data with the SD.sub.2ID of that entry end module and, more importantly, the DD.sub.2ID of the exit end module. Knowing the entry and exit end modules, beneficially removes any translations needed within the transmission network. This allows the data packet to be forwarded across the transmission network directly from the entry end module to any and all intermediate modules, and finally to the targeted exit end module without performing any table lookups as in conventional router techniques. When data is initially transmitted, the SEA external to the transmission network is translated to the SEA of the entry end module, so that any responses can be directed to that SEA which, by the way, has a corresponding D.sub.2ID that is forwarded as an SD.sub.2ID and targeted on responses as a DD.sub.2ID. The same mapping functions of FIG. 7 take place for FIG. 8, with the mapping functions being a relatively small number N. Instead of having to map internal to the transmission network, mapping only takes place based on a limited number of devices N immediately upstream or downstream of an entry or exit end module. By knowing relatively few external modules adjacent the network, mapping can be efficiently performed with little, if any, translations. For example, there may be only one termination device connected to an entry end module, thereby making N equal to one.

Detail Description Paragraph - DETX (41):

[0082] It will be appreciated to those skilled in the art having the benefit of this disclosure that the various embodiments herein are believed to capable of performing fast and efficient transfers across a network which either replaces or forms at least a portion of a conventional communication system. The network may or may not be structured. Regardless of its addressing mechanism or domain, the network advantageously maps to other networks having an addressing domain separate and independent from the network. Various modification and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. The efficient transfers take place, in part, by utilizing a separate set of identification numbers or addresses unique to the network. However, those identification numbers and addresses are transparent to the end-user, and all mapping to those addresses take place at the entry end modules, whereas the added identification numbers and addresses are stripped at the exit end module, with minimal mapping at the network periphery. It is intended that the following claims be interpreted to embrace all such modifications and changes and, accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

Claims Text - CLTX (20):

19. A method for sending a packet of data across a communication network having intermediate modules interposed between an entry end module and an exit end module, comprising: wrapping the packet of data at the entry end module with a source address of the entry end module and a destination address of the exit end module; forwarding the packet of data through the intermediate modules to the exit end module; and mapping the address of the exit end module to an address of a termination device located external to, yet operably coupled with, the communication network, wherein the exit end module has an identification number independent of the termination device, and therefore the exit end module and the termination device are addressable from independent and separate addressing domains or protocols linked only by said mapping.

Claims Text - CLTX (22):

21. The method as recited in claim 19, wherein said mapping comprises: removing the source address and destination address from the packet of data at the exit end module; and while maintaining an upper layer addressing protocol associated with the packet of data, forwarding the packet of data from the exit end module to a termination device addressed by the upper layer addressing protocol.